

HP F8 Architecture

technology brief, 2nd edition



Abstract.....	2
Introduction.....	2
Need for F8 architecture	2
Hot-Plug RAID Memory	5
Memory configuration	5
RAID memory striping	6
Hot-plug memory capabilities	6
Benefits of data protection with RAID.....	7
Error detection and correction	7
Architectural differences from storage subsystem RAID.....	8
F8 crossbar switch	8
Buffer design	8
Multiport design.....	9
Cache coherency filter.....	10
Optimizing cross-bus traffic	11
I/O subsystem.....	11
PCI mode	11
PCI-X mode	12
Intel Xeon processor MP subsystem	12
Hyper-Threading Technology	12
Frequency and full-speed cache	13
Processor and I/O bus design	13
Intel NetBurst architecture.....	13
Conclusion.....	13
Call to action	14

Abstract

HP has developed an 8-way multiprocessing architecture that meets the bandwidth demands of high-end peripherals and the Intel® Xeon™ processor MP. The HP F8 chipset provides key functionality, such as Hot Plug RAID Memory, that was previously unavailable within industry-standard servers. Like redundant array of independent disk technology used in storage subsystems, Hot Plug RAID Memory uses a redundant array of industry-standard DIMMs (RAID) to provide both fault tolerance and the ability to hot replace and hot add memory while the server is operating. The F8 chipset uses a multiport, nonblocking crossbar switch to optimize efficiency and allow simultaneous access to memory, processor, and I/O subsystems. The F8 chipset supports multiple PCI-X bridges and incorporates an embedded HP PCI Hot Plug controller for high availability in the I/O subsystem. The balanced architecture of the F8 chipset delivers superior performance for the most demanding applications, whether they are memory intensive, I/O intensive, or processor intensive.

Introduction

HP has leveraged the experience Compaq gained from the development¹ and use of the Profusion 8-way architecture to design a new 8-way multiprocessing architecture with even higher performance: the F8 architecture. This architecture is based on the Intel® Xeon™ processor MP and is designed to deliver high bandwidth and performance for I/O, processor, and memory subsystems.

The F8 architecture includes HP Hot-Plug RAID Memory—a technology within HP Advanced Memory Protection that is designed for achieving high availability, scalability, and fault tolerance within the memory subsystem. Hot-Plug RAID Memory uses a redundant array of industry-standard DIMMs (RAID) to provide availability and fault tolerance in the memory subsystem, much as redundant array of independent disk (RAID) technology provides availability and fault tolerance in storage subsystems. HP designed the F8 architecture with increased memory bandwidth, a nonblocking crossbar switch that improves bus efficiency, and PCI Hot-Plug and PCI-X capabilities in the I/O subsystem. The ProLiant DL760 G2 and the ProLiant DL740, which use this architecture, vary slightly in implementation. For completed details about these servers, see the [HP website](#).²

Need for F8 architecture

Intel Xeon processors MP operate at speeds greater than 2 GHz and support a bus with four times the bandwidth of the P6 processor bus. (P6 is the family name for Intel processors starting with the Intel Pentium Pro and continuing through the Pentium® III Xeon processor.) Peripherals use high-speed interconnects such as Gigabit Ethernet and Ultra320 SCSI, which operate at bandwidths of 125 MB/s and 320 MB/s, respectively. Clearly, servers need high processor-to-memory bandwidth as well as high I/O-to-memory bandwidth.

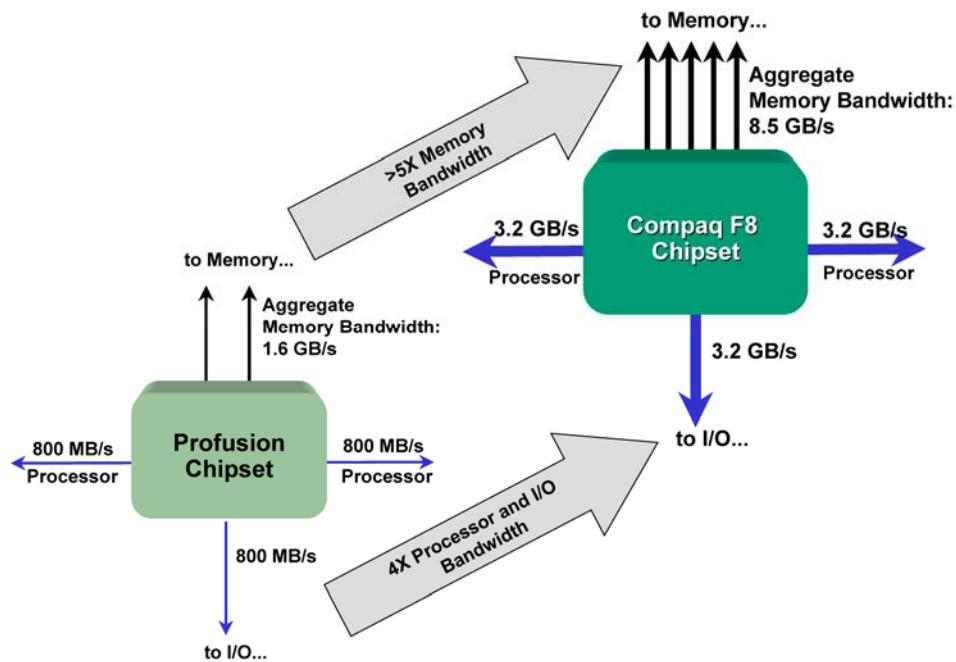
Achieving optimum performance requires a balanced server architecture to ensure that every subsystem—processor, I/O, and memory—has adequate bandwidth. Compaq worked with Corollary to develop the highly successful, balanced architecture in the previous Profusion 8-way chipset. HP has used that experience to design its own 8-way chipset that maximizes bandwidth and performance in all subsystems.

Specifically, the Profusion 8-way architecture had a bus bandwidth of 800 MB/s for the dedicated processor and I/O buses. The F8 architecture is capable of a bandwidth that is four times greater: 3.2 gigabytes per second (GB/s) for each processor bus and for the I/O subsystem (Figure 1).

¹ The Profusion architecture was co-developed by Compaq and Corollary.

² ProLiant DL server information is available at: www.hp.com/servers/dl

Figure 1. Bandwidth comparison between Profusion and F8 architectures



The Profusion architecture ensured fast access to memory by using an aggregate memory bandwidth of 1.6 GB/s. This was enough to balance the maximum bandwidth of the two processor buses in the Profusion architecture. In comparison, the HP F8 architecture ensures even faster memory access by using an aggregate memory bandwidth of 8.5 GB/s, which is 33 percent greater than the bandwidth of the two processor buses combined, and more than five times the memory bandwidth of the previous Profusion architecture.

In the F8 architecture, the total inputs to memory from the two processor buses and the I/O bus provide a cumulative maximum of 9.6 GB/s. The bandwidth to memory is 8.5 GB/s. Thus, the ratio of total inputs to memory to available memory bandwidth (8.5:9.6) approaches an ideal one-to-one ratio, ensuring good scalability for the 8-way multiprocessing architecture (Table 1).

Table 1. Comparison of bandwidth ratios for the Profusion and F8 architectures

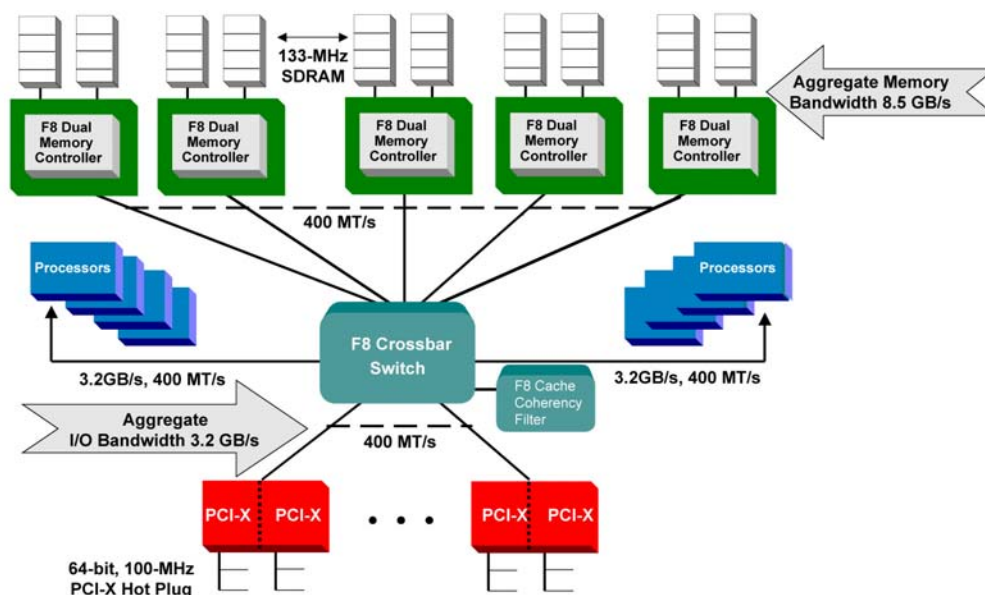
Architecture	Memory bandwidth	Processor buses + I/O bus bandwidth (P1 + P2 + I/O)	Ratio of memory : processor + I/O
Profusion	1.6 GB/s	2.4 GB/s	1.6 : 2.4 (0.67)
F8	8.5 GB/s	9.6 GB/s	8.5 : 9.6 (0.89)

The backbone of the new 8-way architecture is the F8 chipset designed by HP. It includes five memory controllers with patented HP Hot-Plug RAID Memory and a multiported crossbar switch (Figure 2). Product implementations will vary.

The F8 chipset supports:

- An aggregate memory bandwidth of 8.5 GB/s using five separate memory controllers with 400 mega transfers per second³ (MT/s) point-to-point connections. The RAID memory controllers interface with the crossbar switch using a 200-MHz, double-pumped connection to achieve the effective 400 MT/s. Each of the five memory controllers has dual paths into channels of PC100 or PC133 synchronous dynamic random access memory (SDRAM).
- Up to 64 GB of addressable memory.
- Hot -plug RAID Memory, allowing replacement and addition of memory while the server is operating. The RAID design stripes data across multiple memory cartridges while storing parity information in a separate memory cartridge.
- Independent, nonblocking access to memory, processors, and I/O through the multiported crossbar switch. A cache coherency filter reduces the amount of snoop traffic on the processor buses.
- Up to four industry-standard PCI-X bridges, each with an embedded PCI Hot Plug controller. Each of these bridges resides on a 400 MT/s, point-to-point connection, and each bridge can support two PCI-X bus segments operating at speeds up to 100 MHz.
- Up to eight Intel Xeon processors MP. The Intel Xeon processor MP is the multiprocessor version of the seventh-generation IA-32 processor family, designed for high-end workstations and servers.

Figure 2. Block diagram of the F8 chipset architecture



³ Bus speeds are described in mega transfers per second (MT/s). For example, a bus operating at 100 MHz and transferring four data packets on each clock (quad-pumped) would have 400 MT/s. The quad-pumped bus speed at 100 MHz is commonly referred to as 400 MHz rather than 400 MT/s.

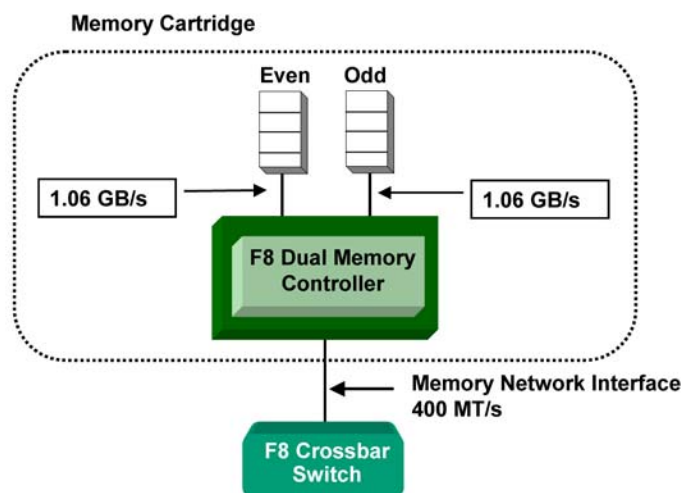
Hot-Plug RAID Memory

Probably the most significant improvement in the F8 architecture is the addition of Hot-Plug RAID Memory, which increases availability, scalability, and fault tolerance in industry-standard servers. The F8 memory controllers provide greatly increased memory bandwidth to handle the system bus speeds, which are four times greater than the P6 bus speeds. The F8 architecture supplies hot-add, hot-replace, and hot-upgrade capabilities. It allows the detection of otherwise undetectable memory errors, which provides a level of data protection far greater than parity or error correcting code (ECC) solutions. HP Hot-Plug RAID Memory enables the memory subsystem to withstand a complete memory device failure and to continue operating normally.

Memory configuration

The F8 chipset uses five memory controllers designed by HP to control five cartridges of industry-standard PC100 or PC133 SDRAM. Within each cartridge, a dual memory controller uses 1.06-GB/s paths into two separate channels of memory (Figure 3). This gives a total bandwidth of 2.12 GB/s within each memory cartridge. External to the memory cartridge, the memory controllers interface with the crossbar switch using a 200-MHz, double-pumped, point-to-point connection. Thus, the memory network interface has an effective data transfer rate of 400 MT/s.

Figure 3. Block diagram showing memory configuration for a single memory cartridge. Each dual memory controller has two independent paths to the two-way interleaved memory channels.




The two memory channels are cache-line interleaved; they share a common address range. As a memory controller performs a write transaction, cache lines with even addresses go to one memory channel and cache lines with odd addresses simultaneously go to the other. Cache-line interleaving is advantageous because memory accesses are typically localized: certain address ranges tend to be accessed more frequently than others, creating "hot spots" in the memory. Interleaving allows the memory controller to split the heavily used locations between the two channels, since roughly half of all accesses will be even and half will be odd.

RAID memory striping

When the memory controller needs to write data to memory, it splits the cache line of data into four blocks. Then each block is written, or striped, across either the even or odd channel of memory in the memory cartridge. A RAID engine in the F8 chipset calculates the Boolean exclusive-OR (XOR) parity information, which is stored on a fifth cartridge dedicated to parity (Figure 4). The four data cartridges and the parity cartridge are each protected by ECC. With the redundant parity data, complete and correct data can be rebuilt from the remaining four cartridges if the data from any DIMM is incorrect or if any cartridge is removed.

Figure 4. Data striping across one of the channels in HP Hot Plug RAID Memory



Because one memory cartridge is dedicated to storing parity information, the architecture has the effective bandwidth of four memory controllers, or 8.5 GB/s (that is, 2.12 GB/s each for four controllers). This is an astounding improvement in performance of the memory interface compared to the 1.6-GB/s aggregate memory bandwidth of the Profusion architecture. HP designed the F8 chipset to take advantage of the faster 400-MT/s memory network interface and to support more memory controllers than the Profusion architecture does.

Each memory controller supports eight DIMMs for a maximum usable memory of 32 GB using 1-GB DIMMs. When using 2-GB DIMMs, the chipset can support up to 64 GB of memory on the four active memory controllers.

It is important to note that HP Hot-Plug RAID Memory has no more performance overhead than standard ECC memory. In Hot-Plug RAID Memory, the RAID engine calculates parity in parallel to the data flow, so no additional data latency is incurred if an error is corrected.

Hot-plug memory capabilities

The redundancy in HP Hot-Plug RAID Memory provides the ability to hot plug memory cartridges, delivering unprecedented levels of memory availability and scalability within industry-standard servers. Hot-Plug RAID Memory enables replacement, addition, and upgrade of DIMMs without shutting down the server.

Hot replace allows a system administrator to replace a failed DIMM while the system is running. Hot replace capability is available in a driverless implementation that requires no support from the operating system. Servers with HP Hot-Plug RAID Memory have hot-replace capability directly out of the box, regardless of the operating system used.

When a hot-replace operation is initiated, the memory controller tells the server to ignore the cartridge of memory where the hot-replace operation will be performed. Until the hot-replace operation is completed, memory transactions use the other four memory cartridges protected by ECC. Thus, the memory subsystem operates in a nonredundant mode like today's ECC memory subsystems. Once the fifth memory cartridge is back online, full redundancy is restored.

When a hot-plug operation is completed, HP Hot-Plug RAID Memory automatically rebuilds the data across all the memory cartridges. Rebuilding data can degrade memory performance briefly. For

example, a rebuild for 4 GB of memory takes less than 30 seconds, a small price to pay to avoid downtime while increasing fault tolerance.

After the RAID engine rebuilds the data, a verify procedure confirms that the rebuild operation was successful. During a verify procedure, every address location in memory is read. Errors found will be reported to the system. If the verify fails, the system continues to operate in non-redundant mode and the new memory will not be brought online until the problem is corrected.

Hot-add and hot-upgrade capabilities allow a user to scale up a computer system as needed by adding or exchanging DIMMs in a memory cartridge while the system is operating. Hot-add and hot-upgrade capabilities require support from the operating system to recognize the additional memory that is available. Several operating systems support hot-add and hot-upgrade, including:

- Windows Server 2003
- SuSE Linux Enterprise Server 7
- Red Hat Enterprise Linux AS 2.1
- SCO UnixWare 7.1.3
- Caldera OpenUnix 8

Benefits of data protection with RAID

Some suppliers of industry-standard servers, including HP, use an alternative data protection method known as distributed ECC to guard against memory device failures. Distributed ECC provides better data protection than standard ECC by distributing bits across multiple DRAM devices. However, if a DRAM device fails, the DIMM must be replaced. Without the redundancy of Hot-Plug RAID Memory, a failed DRAM device results in the need for immediate, unplanned downtime to replace the bad memory DIMM. With HP Hot-Plug RAID Memory, the RAID engine provides redundancy to ensure data protection, and the hot-plug abilities allow a DIMM to be replaced without any downtime.

Error detection and correction

The F8 chipset uses ECC logic in each memory controller to maintain data integrity throughout the memory subsystem. HP has developed an advanced 8-bit ECC algorithm that can reliably detect single-bit, multi-bit, and 4-bit or 8-bit DRAM failures in memory devices. The RAID engine developed by HP corrects these errors (Table 2).

Table 2. Comparison of protection provided by parity checking, ECC, and HP Hot Plug RAID Memory

Error condition	Parity	Standard ECC	Hot Plug RAID Memory
single-bit	detect	correct	correct
double-bit	X	detect	correct
DRAM failure	X	detect	correct
ECC detection fault	X	X	detect

In a memory read transaction, every block of data simultaneously travels through the ECC logic and the RAID parity engine. The ECC logic determines whether the data is good or bad. If the data is bad, the chipset uses the regenerated data from the RAID engine. Thus, the error detected by the ECC is eliminated and only good data is transmitted.

If the ECC logic sends a signal that the data is good, then this data is compared with the regenerated data from the RAID engine. If the two blocks of data are not identical, an error undetectable by ECC has occurred. While such an occurrence would be rare, an ECC-only system would be unable to detect such failures and could pass along corrupt data as if it were good.

With HP Hot-Plug RAID Memory, when an error undetectable by ECC occurs, the data comparison fails and the memory controller initiates a nonmaskable interrupt (NMI), preventing transmission of corrupt data. This feature makes HP Hot-Plug RAID Memory virtually immune to data corruption.

Architectural differences from storage subsystem RAID

The technology used in HP Hot-Plug RAID Memory is conceptually similar to RAID technology that provides fault tolerance and high availability in storage subsystems for servers. However, there are some key performance and implementation differences between Hot-Plug RAID Memory and typical storage subsystem RAID.

Hot-Plug RAID Memory does not have the mechanical delays of seek time and rotational latency associated with hard disk drive arrays. Storage subsystem arrays use a single bus to write the stripes sequentially across multiple drives. In contrast, HP Hot-Plug RAID Memory uses parallel point-to-point connections so that data is written simultaneously across multiple memory cartridges.

Also, HP Hot-Plug RAID Memory eliminates the write bottleneck associated with typical storage subsystem RAID implementations. In a storage array, the RAID controller generally performs a read operation of existing parity before a write operation can be completed. If a dedicated parity drive is being used, a bottleneck occurs. But because HP Hot-Plug RAID Memory operates on an entire cache line of data, there is no need to read existing parity before a write operation, thus eliminating this performance bottleneck.

When a traditional striped RAID storage subsystem rebuilds data, there is no data protection should another drive fail. However, the F8 chipset operates in a typical (nonredundant) ECC mode while data is being rebuilt. As a result, even if a secondary memory failure occurs during a rebuild operation, the data is protected by ECC.

F8 crossbar switch

One of the key advantages that the Profusion architecture has over other 8-way designs is its use of a nonblocking, multiported crossbar switch. This switch allows simultaneous communication among the processors, I/O, and memory. The F8 architecture also uses a nonblocking, multiported crossbar switch that provides even higher performance than the Profusion crossbar switch and accommodates increased processor speeds and peripheral bandwidths. The F8 chipset also includes a cache coherency filter, or cache accelerator, similar to that in the Profusion architecture. The cache coherency filter removes (or filters) unnecessary snoop cycles on the processor buses.

HP engineers designed the F8 crossbar switch to increase bus efficiency far beyond that of the Profusion crossbar switch. The design includes:


- **Larger and reorganized buffers.** The F8 crossbar switch can hold 128 cache lines, twice the number that the Profusion chipset can hold in its buffers.
- **More ports.** The F8 crossbar switch has thirteen read and four write ports, compared with five read and five write ports used in the Profusion chipset. This increases the number of transactions that can run concurrently.
- **Optimized cross-bus traffic through a patent-pending algorithm.** Optimizing the cross-bus traffic significantly enhances the ability to scale beyond 4-way multiprocessing.

Buffer design

The Profusion chipset uses a single centralized buffer, or queue, for storing data requests. In certain cases, a processor on one bus could request the same address as a processor on the other bus, resulting in the need to arbitrate for which request could be granted first. One of the requests has to go through a retry process, using up additional bandwidth on the processor bus.

In the F8 architecture, the crossbar switch (Figure 5) contains a separate buffer for each of the processor buses, the I/O subsystem, and the memory subsystem. The buffers in the crossbar switch are distributed so that the data is stored closest to where it enters the application-specific integrated circuit (ASIC).

Figure 5. The F8 crossbar switch uses distributed buffers and multiple read and write ports.



With the F8 crossbar switch, the request is logged into the appropriate buffer, and then each request is processed in a fair-share algorithm. The distributed buffer design and the increased buffer sizes reduce the amount of arbitration and the number of retry cycles required when processors request information, allowing the processors to do more useful work.

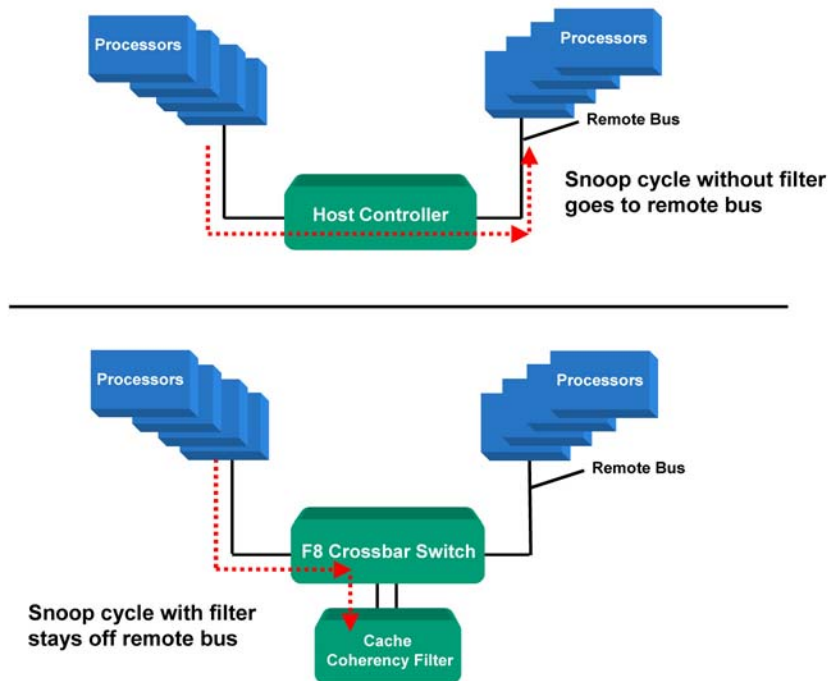
Multiport design

The F8 crossbar switch contains four write ports and thirteen read ports (Figure 5) and allows simultaneous data transfer on any of those ports. By comparison, the Profusion chipset has five read ports and five write ports. Despite having fewer write ports than Profusion chipset, the F8 crossbar switch significantly improves performance because its port to main memory is extremely wide, with a bandwidth more than five times greater than that of the Profusion chipset.

Cache coherency filter

One of the challenges of designing an efficient multiprocessing architecture is to maintain a consistent view of memory by all the processors and the I/O subsystem. This is typically referred to as maintaining cache coherency. Because data is shared among several level two (L2) caches on the processors, it is possible that data referred to by two different caches could be inconsistent. In a multiprocessing server with dual processor buses, a memory transaction from one processor bus has to look at, or snoop, the remote processor bus to make sure that only the most recent data is in use. Every snoop cycle consumes bandwidth on the remote processor bus and diminishes the performance of the system (Figure 6).

Figure 6. Comparison of a snoop cycle with and without a cache coherency filter



The F8 chipset uses a cache coherency filter to reduce the number of snoop cycles on the remote processor bus. The cache coherency filter is also known as a cache accelerator. It holds the addresses of data stored in all of the L2 processor caches, as well as information about the state of the data. For example, the state information may describe whether the data is owned by a particular L2 cache or shared between multiple caches.

The cache coherency filter also acts as a filter for the I/O bus, keeping track of which cache lines are owned on the I/O bus for the PCI devices. When a processor requests a cache line, the crossbar switch snoops the I/O filter to determine if that cache line resides in one of the PCI bridges on the I/O bus. If the cache line is not present in one of the bridges, then no transaction is run on the I/O bus. This reduces snoop traffic on the I/O bus whenever a processor requests data.

Optimizing cross-bus traffic

The F8 chipset alleviates some inefficiency that the Profusion chipset has when snoop traffic must cross to the remote processor bus. When a processor requests data, the Profusion chipset checks the cache coherency filter to determine the specific location of the data it needs. If the data is located in an L2 cache on the remote bus, the chipset snoops the remote bus to obtain the data, causing cross-bus traffic. In the Profusion chipset, a read request that requires a snoop cycle on the remote bus is automatically deferred,⁴ causing a reply to be sent at a later time. This situation generates two cycles on the processor bus for every single read request.

The F8 chipset optimizes cross-bus traffic by incorporating a patent-pending Guaranteed Snoop Access algorithm. The algorithm defers fewer requests than Profusion does, thus reducing the amount of traffic on the processor bus. The F8 chipset defers cycles only when necessary to prevent a livelock situation⁵, yet maintains the order and coherency of the requests. Through the Guaranteed Snoop Access algorithm, HP designers have significantly optimized the flow of cross-bus traffic and thus enhanced the scalability of the F8 architecture.

I/O subsystem

HP is a leading technology innovator of industry-standard I/O subsystems, as evidenced by its development of PCI Hot-Plug technology, the I/O controller for the Profusion chipset, and co-development of the latest enhancement to the PCI bus: PCI-X technology.

HP has used this expertise to help a chipset vendor develop an industry-standard PCI-X bridge that provides a high-performance data path between the F8 chipset and peripheral devices. HP designed the F8 chipset to support up to four of these industry-standard PCI-X bridges using a 200-MHz, double-pumped, point-to-point connection that results in an effective data transfer of 400 MT/s. The point-to-point connection is source synchronous, which means that the clock signal travels with the data signal. Because the clock signal and the data travel together, the risks of signal degradation are minimized and the source signal is always synchronized with the receiver to provide more effective data transmission.

Each PCI-X bridge supports two 64-bit PCI-X bus segments. Each of the eight bus segments can be independently configured to run either in PCI mode operating at 33 or 66 MHz or in PCI-X mode operating at 66 or 100 MHz. Both modes support PCI Hot Plug using an integrated controller developed and licensed by HP.

PCI mode

The PCI-X bridge supports delayed PCI transactions, an important feature that improves bus performance. All reads to main memory are completed as delayed transactions when the PCI-X bridge operates in PCI mode. The device that initiates the transaction polls the PCI-X bridge to determine if the requested data is cached there, rather than holding the bus while waiting for the data. This polling allows other devices to use the bus while the transaction is completed.

The PCI-X bridge includes prefetch buffers to make it a caching device. Each buffer can hold multiple cache lines. These buffers have been sized to provide optimal performance at a reasonable and cost-effective silicon die size. Because of the delayed transaction support, the PCI-X bridge can get data for multiple PCI devices concurrently.

⁴ •A deferred request is split into two transactions so that the processor makes a read request and gets off the bus. Then a reply is sent when the data is available.

⁵ Livelock: When two processes continuously change their state in response to changes in the other process without doing any useful work. (The Free On-line Dictionary of Computing, <http://foldoc.doc.ic.ac.uk/>, Editor Denis Howe)

PCI-X mode

The F8 architecture incorporates PCI-X technology to significantly expand the I/O performance. PCI-X technology, developed by Compaq, Hewlett-Packard, and IBM, is an evolutionary I/O upgrade to conventional PCI technology. PCI-X enables the design of I/O subsystems and peripheral devices that can operate at bus frequencies greater than 66 MHz using a 64-bit bus width. The PCI-X bridge designed for the F8 architecture runs at 66 or 100 MHz, allowing flexibility for system architects and supporting multiple devices for end users.

PCI-X improves performance over conventional PCI as a result of two primary differences: higher clock frequencies made possible by a register-to-register protocol and new protocol enhancements, such as split transactions, to make the bus more efficient. The register-to-register protocol eases the timing constraints by allowing an entire clock cycle for the decode logic to occur. With the timing constraints reduced, it is much easier to design adapters and systems to operate at frequencies greater than 66 MHz.

In PCI-X mode, read operations to main memory are completed as split transactions rather than as delayed transactions. A split transaction enables more efficient use of the bus because it eliminates polling. With a delayed transaction in conventional PCI protocol, the device requesting data must poll the target to determine when the request has been completed and the data is available. With a split transaction as supported in PCI-X, the device requesting the data sends a signal to the target. The target device informs the requester that it has accepted the request. The requester is free to process other information until the target device sends the data to the requester.

The F8 architecture includes two optional features from the PCI-X specification to enhance performance even more: the “don’t-snoop” bit and relaxed ordering. When the “don’t snoop” bit is set during a PCI-X transaction, an I/O request will not snoop the L2 caches on the processor bus. Thus, an I/O request will go directly to main memory, eliminating a snoop cycle on the processor bus.

With conventional PCI bridge designs, the bridge handles requests from multiple PCI devices in the order in which they are received. The PCI-X protocol includes an optional relaxed ordering bit. If the device driver or controlling software sets this bit, the PCI-X bridge permits a transaction to pass previously posted transactions from other devices. The bridge can rearrange the transactions in the most efficient manner, depending on which PCI device or system memory port is available.

Intel Xeon processor MP subsystem

The Intel Xeon processor MP is the multiprocessing version of the seventh-generation IA-32 processors.⁶ The Intel Xeon processor MP is based on the Intel NetBurst® architecture and is designed for performance in high-end x86 workstations and servers. The seventh-generation architecture is significantly different than the architecture of the Intel P6 family, which began with the Pentium Pro and extended through the Pentium III Xeon processors.

Hyper-Threading Technology

The Intel Xeon processor MP uses Intel Hyper-Threading technology that improves processor utilization to meet the needs of large, memory-intensive server applications. Hyper-Threading technology enables one physical processor to execute two separate threads at the same time. To achieve this, Intel designed the Xeon processor MP with the usual processor core, but with two Architectural State devices (logical processors). Each Architectural State tracks the flow of a thread being executed by core resources. Both logical processors inside the physical processor share all the internal caches and other physical execution resources. An application or operating system can submit threads to two different logical processors just as it would in a traditional multiprocessor system. The execution core

⁶ More detailed information about the Xeon MP processor is available in the Technology Brief entitled The Intel® processor roadmap for industry-standard servers, <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00164255/c00164255.pdf>

processes instructions in an order determined by dependencies in the data and availability. Therefore, the processor is allowed to execute instructions in the order that will yield the best overall performance.

For more information, see the HP technology brief⁷ entitled "[Intel® Hyper-Threading Technology](#)."

Frequency and full-speed cache

The Intel Xeon processors MP used in the ProLiant DL760 G2 and ProLiant DL740 servers are available with operating frequencies of up to 3.0 GHz (using the 130 nm processing technology). The Intel Xeon processor MP includes an L2 cache located on the same die as the processor logic, giving high bandwidth and low latency on a full-speed backside bus. The full-speed backside bus enables efficient access to the most frequently used data. The Intel Xeon processor MP also includes an integrated level three (L3) cache on the die with size options of 1, 2, or 4 MB.

Processor and I/O bus design

The 64-bit system bus for the Intel Xeon processor MP uses a similar protocol and cache coherency design as the P6 bus. The bus operates at 100 MHz using a quad-pumped data rate. The quad-data-rate bus uses four separate clocks, or strobes, to allow data transfer four times within a single clock cycle; therefore it provides an effective data transfer frequency of 400 MT/s and a maximum theoretical bandwidth of 3.2 GB/s.

Intel NetBurst architecture

The NetBurst architecture uses a hyper-pipeline, a 20-stage branch prediction pipeline that can contain more than 100 instructions at once and can handle up to 48 loads and stores concurrently. Specific to this NetBurst design is an improved branch-prediction algorithm to mitigate effects of branch mispredicts on the long pipeline. The NetBurst architecture also includes:

- Support for Streaming SIMD extension 2 (SSE2), to manage floating point, application, and multimedia performance.
- Deeper instruction window for out-of-order, speculative execution and improved branch prediction over the P6 dynamic execution core
- A double-data rate arithmetic logic unit that is clocked at twice the speed of the processor.
- Execution trace cache to store pre-decoded micro-operations.

Conclusion

The F8 chipset delivers bandwidth four to five times greater than that in the previous 8-way Profusion architecture. It is capable of providing the performance and uptime required to meet the demands of enterprise server consolidation, database, and data mining/warehousing applications. Its nonblocking crossbar switch allows direct point-to-point access to all system resources: processors, memory, and I/O. The balanced architecture of the F8 chipset delivers superior performance for the most demanding applications, regardless of whether these applications are processor intensive, memory intensive, or I/O intensive. Perhaps most importantly, HP has developed the new capability of Hot-Plug RAID Memory to provide an unprecedented level of fault tolerance, scalability, and availability while using industry-standard DIMMs.

⁷ The technology brief is available on the HP website at <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00257074/c00257074.pdf>

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

TC050606TB, 06/2005

